arXiv:1606.01487v1 [stat.ML] 5 Jun 2016

Bounds for Vector-Valued Function Estimation

Andreas Maurer
Adalbertstrasse 55, D-80799 Munchen, Germany
Email: *am@andreas-maurer.eu*

Massimiliano Pontil
Istituto Italiano di Tecnologia, 16163 Genoa, Italy
Email: *massimiliano.pontil@iit.it*
and
University College London
Department of Computer Science, London WC1E 6BT, UK

June 7, 2016

Abstract

We present a framework to derive risk bounds for vector-valued learning with a broad class of feature maps and loss functions. Multi-task learning and one-vs-all multi-category learning are treated as examples. We discuss in detail vector-valued functions with one hidden layer, and demonstrate that the conditions under which shared representations are beneficial for multi-task learning are equally applicable to multi-category learning.

1 Introduction

The main focus of this paper is to study statistical bounds for (shared) representation learning under a general class of feature maps and loss functions. This study is motivated by the development of data-dependent generalization bounds for multi-category learning with T classes, and for multi-task learning with T tasks. We show that both problems can be treated in parallel under a unified framework.

We give bounds on the Rademacher complexity of composite vector-valued function classes

$$\mathcal{F} \circ \mathcal{G} = \left\{ x \in H \mapsto f\left(g\left(x\right)\right) \in \mathbb{R}^T : f \in \mathcal{F}, g \in \mathcal{G} \right\},$$

where the input space H is a finite or infinite dimensional Hilbert space, $\mathcal{G}$ is a class of functions (or feature-maps or representations) $g : H \to \mathbb{R}^K$, and $\mathcal{F}$ is a class of output functions $f : \mathbb{R}^K \to \mathbb{R}^T$.

Functions in $\mathcal{F} \circ \mathcal{G}$ are chosen on the basis of a finite number N of independent observations and we are interested in uniformly bounding the incurred estimation errors in terms of the parameters T, K and N, or alternatively $n = N/T$, the number of observations per output unit.

There are two main contributions of this work:

- We provide a common method to derive data dependent bounds for multi-task and multi-category learning in terms of the complexity of general vector-valued function classes. In passing we improve on a recent result in [15] on multi-category learning. Our framework is also general enough to be applied to hybrid coding schemes for multi-category classification such as 1-vs-1 pairwise classification.
- We apply this method to a large class of vector-valued functions with shared feature maps to demonstrate that the conditions under which shared representations are beneficial for multi-task learning are equally applicable to multi-category learning.

Our principal finding is a data-dependent generalization bound, whose dominant terms have the form

$$O\left(\theta\sqrt{\frac{\mathrm{tr}(\hat{C})}{nT}}\right) + O\left(\theta\sqrt{\frac{\lambda_{\max}(\hat{C})}{n}}\right),$$

where $\hat{C}$ is the empirical covariance operator (see below). When testing multi-task learning we are always told which task we are testing and thus the relevant component of our vector-valued hypothesis. In the one-vs-all multi-category setting we of course withhold the identity of the correct class and thus also of the relevant component. This simple fact is reflected in the presence of the factor θ, which is one for multi-task learning and $\sqrt{T}$ for multi-category learning.

Bounds of this form are given for a large class of neural networks with one hidden layer and rather general nonlinear activation functions, which may involve inter-unit couplings or intermediate maps to infinite-dimensional spaces. A similar bound also holds for linear classes with trace-norm constraints, which can also be interpreted as composite classes, see e.g. [26].

As T increases the second term dominates the above expression. This term however depends only on the largest eigenvalue, instead of the trace, of the empirical covariance. If T is large and the data is high-dimensional the intermediate representation can therefore give a considerable advantage. This has been established for multi-task learning in several works and, as we show here, holds equally for multi-category learning, in agreement with previous empirical studies of the benefit of trace-norm regularization in multi-category learning [1].

In Section 2 we explain how the complexities of multi-category and multi-task learning can be reduced to the complexities of vector-valued function classes and bounded by a common expression. We briefly discuss independent and linear classes in Section 3.1 and 3.2. Then in Section 3.3, we present our principal result on nonlinear composite classes. The appendix contains statements and proofs of our results in their most general form.

1.1 Previous Work

Bounds for multi-layered networks are given in the now classical work [3] in terms of covering numbers. More recently there are bounds using Rademacher averages [24]. These works mainly consider scalar outputs and ignore the regularizing effects of intermediate representations.

Early work to consider the potential benefits of shared representations was in the setting of multi-task learning and learning to learn [5]. Subsequent work has focused more on learning bounds for linear feature learning [7, 16]. Recently [20] presented a general bound for multi-task representation learning. Although there has been substantial work on the statistical analysis of learning shared representations for multi-task learning, less has been done for multi-category learning. This is in contrast with the large body of empirical work on deep networks, which are often trained with a multi-class loss [9], such as the soft max or multi-class hinge loss. In this work we close this gap.

2 Multi-Category and Multi-Task Learning

We extend the notion of Rademacher complexity to the vector-valued setting.

Definition 2.1 *Let $T, N \in \mathbb{N}$, let $\mathcal{X}$ be any set, $\mathcal{F}$ a class of functions $f : \mathcal{X} \to \mathbb{R}^T$, $\mathbf{x} = (x_1, \ldots, x_N) \in \mathcal{X}^N$, and let $I : \{1, \ldots, T\} \to 2^{\{1,\ldots,N\}}$ be a function which assigns to every $t \in \{1, \ldots, T\}$ a subset $I_t \subset \{1, \ldots, N\}$. We define*

$$R_I(\mathcal{F}, \mathbf{x}) = \frac{1}{N}\mathbb{E}\sup_{f\in\mathcal{F}}\sum_{t=1}^{T}\sum_{i\in I_t}\epsilon_{ti} f_t(x_i),$$

where the ϵ_{ti} are doubly indexed, independent Rademacher variables (uniformly distributed on $\{-1, 1\}$).

In this section we show that the estimation problem for both multi-category and multi-task learning can be reduced to the problem of bounding $R_I(\mathcal{F}, \mathbf{x})$ for appropriate choices of the function I.

2.1 Multi-Category Learning

Let $C \in \mathbb{N}$ be the number of categories. There is an unknown distribution μ on $H \times \{1, \ldots, C\}$, a classification rule $cl : \mathbb{R}^T \to \{1, \ldots, C\}$, and for each label $y \in \{1, \ldots, C\}$ a surrogate loss function $\ell_y : \mathbb{R}^T \to \mathbb{R}_+$. The loss function ℓ_y is designed so as to upper bound or approximate the indicator function of the set $\{z \in \mathbb{R}^T : cl(z) \neq y\}$. Here we consider the simple case, where $T = C$. For the construction of appropriate loss functions see [8, 15, 23]. These loss functions are Lipschitz on $\mathbb{R}^T$ relative to the Euclidean norm, with some Lipschitz constant L_{mc}, often interpretable as an inverse margin.

Given a class $\mathcal{F}$ of functions $f : H \to \mathbb{R}^T$ we want to find $f \in \mathcal{F}$ so as to approximately minimize the surrogate risk

$$\mathbb{E}_{(x,y)\sim\mu} \ell_y \left(f \left(x \right) \right).$$

Since we do not know the distribution μ, this is done on the basis of a sample of $N = nT$ observations $(\mathbf{x}, \mathbf{y}) = \left(\left(x_1, y_1 \right), \ldots, \left(x_N, y_N \right) \right) \in \left(H \times \{1, \ldots, C\} \right)^N$, drawn i.i.d. from the distribution μ. We then solve the problem

$$\hat{f} = \arg\min_{f \in \mathcal{F}} \frac{1}{N} \sum_{i=1}^{N} \ell_{y_i} \left(f \left(x_i \right) \right).$$

To give a performance guarantee for $\hat{f}$ we would like to know how far the empirical minimum above is from the true surrogate risk of $\hat{f}$. This difference is upper bounded by

$$\sup_{f \in \mathcal{F}} \left[\mathbb{E}_{(x,y)\sim\mu} \ell_y \left(f \left(x \right) \right) - \frac{1}{N} \sum_{i=1}^{N} \ell_{y_i} \left(f \left(x_i \right) \right) \right].$$

It is by now well known (see e.g. [4]) that the above expression has, with high probability in the sample, a bound, whose dominant term is given by

$$\frac{2}{N} \, \mathbb{E} \sup_{f \in \mathcal{F}} \sum_{i=1}^{N} \epsilon_i \ell_{y_i} \left(f \left(x_i \right) \right), \tag{1}$$

where the ϵ_i are independent Rademacher (uniform $\{-1, 1\}$-distributed) variables. We now apply the following result [21, Corollary 6].

Theorem 2.2 *Let $\mathcal{X}$ be any set, $(x_1, \ldots, x_n) \in \mathcal{X}^n$, let $\mathcal{F}$ be a class of functions $f : \mathcal{X} \to \mathbb{R}^T$ and let $h_i : \mathbb{R}^T \to \mathbb{R}$ have Lipschitz norm bounded by L. Then*

$$\mathbb{E} \sup_{f \in \mathcal{F}} \sum_{i=1}^{n} \epsilon_i h_i \left(f \left(x_i \right) \right) \leq \sqrt{2} L \mathbb{E} \sup_{f \in \mathcal{F}} \sum_{t,i} \epsilon_{ti} f_t \left(x_i \right),$$

where ϵ_{ti} is an independent doubly indexed Rademacher sequence and f_t is the t-th component of f.

Using this theorem and the Lipschitz property of the loss functions ℓ_{y_i}, we upper bound (1) by

$$\frac{2\sqrt{2}}{N} L_{\mathrm{mc}} \, \mathbb{E} \sup_{f \in \mathcal{F}} \sum_{t=1}^{T} \sum_{i=1}^{N} \epsilon_{ti} f_t \left(x_i \right). \tag{2}$$

A similar argument can be based on Slepian's inequality with a passage to Gaussian complexities [15]. In this case the ϵ_{ti} have to be replaced by independent standard normal variables γ_{ti}, and $\sqrt{2}$ replaced by $\sqrt{\pi/2}$. The approach chosen here is simpler and allows us to improve some results of [15] in the linear case. For our final result (Theorem 3.3 below) however we also need Gaussian complexities.

We define $I^{\mathrm{mc}} : \{1, \ldots, T\} \to 2^{\{1, \ldots, N\}}$ by $I_t^{\mathrm{mc}} = \{1, \ldots, N\}$ for all t. With Definition 2.1 the quantity (2) then becomes

$$2\sqrt{2} L_{\mathrm{mc}} R_{I^{\mathrm{mc}}} \left(\mathcal{F}, \mathbf{x} \right). \tag{3}$$

2.2 Multi-Task Learning

In this setting there is an output space $\mathcal{Y}$, and for each task $t \in \{1, \ldots, T\}$ a distribution μ_t on $H \times \mathcal{Y}$ and a loss function $\ell_t : \mathbb{R} \times \mathcal{Y} \to [0, 1]$, which is assumed to be Lipschitz with constant at most L_{mt} in the first argument for every value of the second. Given a class $\mathcal{F}$ of functions $f : H \to \mathbb{R}^T$ we want to find $f \in \mathcal{F}$ so as to approximately minimize the task-average risk

$$\frac{1}{T} \sum_{t=1}^{T} \mathbb{E}_{(x,y)\sim\mu_t} \ell_t \left(f_t \left(x \right), y \right),$$

where f_t is the t-th component of the function f. For each task t there is a sample $(\mathbf{x}_t, \mathbf{y}_t) = ((x_{t1}, y_{t1}), \ldots, (x_{tn}, y_{tn}))$ drawn i.i.d. from μ_t. One solves the problem

$$\hat{f} = \arg\min_{f\in\mathcal{F}} \frac{1}{nT} \sum_{t=1}^{T} \sum_{i=1}^{n} \ell_t \left(f_t \left(x_{ti} \right), y_i \right).$$

As before we are interested in the supremum of the estimation difference

$$\sup_{f\in\mathcal{F}} \frac{1}{T} \sum_{t=1}^{T} \left[\mathbb{E}_{(x,y)\sim\mu_t} \ell_t \left(f_t \left(x \right), y \right) - \frac{1}{n} \sum_{i=1}^{n} \ell_t \left(f_t \left(x_{ti} \right), y_i \right) \right].$$

As shown in [2] or [16] there is again a high probability bound, whose dominant term is given by the vector-valued Rademacher complexity

$$\frac{2}{nT} \mathbb{E} \sup_{f\in\mathcal{F}} \sum_{t=1}^{T} \sum_{i=1}^{n} \epsilon_{ti} \ell_t \left(f_t \left(x_{ti} \right), y_i \right) \leq \frac{2}{nT} L_{\mathrm{mt}} \, \mathbb{E} \sup_{f\in\mathcal{F}} \sum_{t=1}^{T} \sum_{i=1}^{n} \epsilon_{ti} f_t \left(x_{ti} \right),$$

where we eliminated the Lipschitz functions with a standard contraction inequality as in [22]. We now collect all the tasks input samples $\mathbf{x}_t$ in a big sample $\mathbf{x} = (x_1, \ldots, x_N) \in H^N$ with $N = nT$, and define $I^{\mathrm{mt}} : \{1, \ldots, T\} \to 2^{\{1,\ldots,N\}}$ so that I_t^{mt} is the set of all indices of the examples for task t. Thus $\mathbf{x}_t = (x_i)_{i\in I_t}$ and $n = \left| I_t^{\mathrm{mt}} \right|$. The right hand side above again becomes

$$2 L_{\mathrm{mt}} R_{I^{\mathrm{mt}}} \left(\mathcal{F}, \mathbf{x} \right). \tag{4}$$

2.3 A Common Expression to Bound

Comparing (3) and (4) we can summarize: Let $\mathcal{F}$ be a class of functions with values in $\mathbb{R}^T$. The empirical Rademacher complexity of $\mathcal{F}$ as used in multi-category learning and the empirical Rademacher complexity of $\mathcal{F}$ as used in multi-task learning are up to (Lipschitz-) constants, bounded by $R_I \left(\mathcal{F}, \mathbf{x} \right)$, where the function I is either I^{mc} in the multi-category case or I^{mt} in the multi-task case and $I_t^{\mathrm{mc}} = \{1, \ldots, N\}$ while $I_t^{\mathrm{mt}} \subseteq \{1, \ldots, N\}$ is the set of indices of examples for task t.

With appropriate definitions of the function I, bounds on $R_I \left(\mathcal{F}, \mathbf{x} \right)$ also lead to learning bounds in hybrid situations where there are several multi-category tasks, potentially with classes occurring

in more than one task. In the case of 1-vs-1 voting schemes $T = C(C-1)/2$, so there is a component for every unordered pair of distinct classes (c_1, c_2). Then we define a $I_{(c_1,c_2)}$ to be the set of indices of all examples for the classes c_1 and c_2.

In general I_t should be the set indices of those examples, which occur as arguments of f_t in the expression of the empirical error. For reasons of space however we will stay with the cases of multi-task and 1-vs-all multi-category learning as explained above. We refer to the appendix for the most general statements of our results.

To lighten notation we write $R_{I^{\mathrm{mc}}} = R_{\mathrm{mc}}$ and $R_{I^{\mathrm{mt}}} = R_{\mathrm{mt}}$. We also use the notation R_α, where the variable α can be either "mc" or "mt". It will also be useful to observe that for $(a_1, \ldots, a_N) \in \mathbb{R}^N$

$$\sum_{t=1}^{T} \sum_{i \in I_t^\alpha} a_i = \theta_\alpha^2 \sum_{i=1}^{N} a_i,$$

where $\theta_{\mathrm{mc}} = \sqrt{T}$ and $\theta_{\mathrm{mt}} = 1$.

3 Specific Bounds

We show how the quantity $R_\alpha(\mathcal{F}, \mathbf{x})$ may be bounded, first by a simple and general method of reduction to the Rademacher complexities of scalar function classes, then for certain linear classes, and finally we state and prove our main results for composite classes.

3.1 Component Classes and Independent Learning

Given a class $\mathcal{F}$ of functions with values in $\mathbb{R}^T$ we can define for each $t \in \{1, \ldots, T\}$ the scalar valued component class $\mathcal{F}_t = \{f_t : f \in \mathcal{F}\}$. By bringing the supremum inside the first sum in (4) we obtain the bound

$$R_I(\mathcal{F}, \mathbf{x}) \leq \frac{1}{N} \sum_{t=1}^{T} \mathbb{E} \sup_{f \in \mathcal{F}_t} \sum_{i \in I_t} \epsilon_i f(x_i),$$

which is just a sum of standard, scalar case, empirical Rademacher averages.

In the case of independent learning the components of the members of $\mathcal{F}$ are chosen independently, so that

$$\mathcal{F} = \prod_t \mathcal{F}_t = \{(f_1, \ldots, f_T) : \forall t, f_t \in \mathcal{F}_t\},$$

and the above bound becomes an identity and unimprovable. In most cases $\mathbb{E} \sup_{f \in \mathcal{F}_t} \sum_{i \in I_t} \epsilon_i f(x_i)$ is of the order $\sqrt{|I_t|}$ so the above implies a bound of the order $\theta_\alpha / \sqrt{N}$.

3.2 Linear Classes

Before proceeding we require some more notation. Given a sequence of input vectors, $(x_1, \ldots, x_N) \in H^N$ we define the empirical covariance operator $\hat{C}$ by

$$\langle \hat{C}\, v, w \rangle = \frac{1}{N} \sum_{i=1}^{N} \langle v, x_i \rangle \langle x_i, w \rangle \text{ for every } v, w \in H\,.$$

Furthermore, given a function $I : \{1, \ldots, T\} \to 2^{\{1,\ldots,N\}}$, we define the empirical covariance operator $\hat{C}_t$ by

$$\langle \hat{C}_t v, w \rangle = \frac{1}{|I_t|} \sum_{i \in I_t} \langle v, x_i \rangle \langle x_i, w \rangle.$$

We consider linear transformations $W : H \to \mathbb{R}^T$ of the form

$$x \mapsto (\langle w_1, x \rangle, \ldots, \langle w_T, x \rangle)$$

with weight-vectors $w_t \in H$. Corresponding function classes will be defined by constraints on the norms of such transformations. We use the mixed $(2,p)$-norms which are defined as

$$\|W\|_{2,p} = \big\|(\|w_1\|, \ldots, \|w_T\|)\big\|_p$$

and the trace norm $\|\cdot\|_{\mathrm{tr}} = \mathrm{tr}\big(\sqrt{W^* W}\big)$. The norm $\|\cdot\|_{2,2}$ is also known as the Hilbert-Schmidt norm or, for finite-dimensional H, as the Frobenius norm $\|W\|_{2,2} = \sqrt{\sum_t \|w_t\|^2}$. For $B > 0$ we consider the classes,

$$\mathcal{W}_{2,p} = \Big\{ W : \|W\|_2 \leq B T^{1/p} \Big\}$$

and

$$\mathcal{W}_{\mathrm{tr}} = \Big\{ W : \|W\|_{\mathrm{tr}} \leq B\sqrt{T} \Big\}\,.$$

The class $\mathcal{W}_{\mathrm{tr}}$ can be defined alternatively as $\mathcal{W}_{\mathrm{tr}} = \{VW : W \in \mathcal{W}, V \in \mathcal{V}\}$, where $\mathcal{W} = \{W : H \to \mathbb{R}^T, \|W\|_{2,2} \leq 1\}$ and $\mathcal{V} = \{V : \mathbb{R}^T \to \mathbb{R}^T, \|V\|_{2,2} \leq B\sqrt{T}\}$, see for example [26] and references therein. This exhibits $\mathcal{W}_{\mathrm{tr}}$ as a composite vector-valued function class.

The factor $T^{1/p}$ in the definition of $\mathcal{W}_{2,p}$ is essential when discussing the dependence on T. If it were absent then by Jensen's inequality the average norm allowed to the weight vectors would be bounded by $B/T^{1/p}$, so the class is regularized to death as T increases. This applies in particular to the case of multi-category learning, where each component needs to be able to win over all the others by some margin. The same argument applies to the $\sqrt{T}$ in the constraint of the trace-norm class. In this sense it is not quite correct to speak of rates in T if the constraint on the norm is held constant as in [15].

For simplicity we assume that $\|x_i\| = 1$ for all i (as with a Gaussian RBF-kernel) for the rest of this subsection. Note that this implies $\mathrm{tr}(\hat{C}) = \mathrm{tr}(\hat{C}_t) = 1$. We also consider only the cases of multi-category and multi-task learning. Statements and proofs for general index sets I_t and general values of the $\|x_i\|$ are given in the appendix. We first give some lower and upper bounds for $\mathcal{W}_{2,\infty}$ and $\mathcal{W}_{2,p}$.

Theorem 3.1 *For $p \in [2, \infty]$*

$$B\ \theta_\alpha \sqrt{\frac{1}{2n}} \leq R_\alpha\left(\mathcal{W}_{2,\infty}, \mathbf{x}\right) \leq R_\alpha\left(\mathcal{W}_{2,p}, \mathbf{x}\right) \leq B\ \theta_\alpha \sqrt{\frac{1}{n}}$$

and for $p \in [1, 2]$ and $1/p + 1/q = 1$

$$R_\alpha\left(\mathcal{W}_{2,2}, \mathbf{x}\right) \leq R_\alpha\left(\mathcal{W}_{2,p}, \mathbf{x}\right) \leq 2^{1/q} B\ \theta_\alpha \sqrt{\frac{q}{n}}.$$

The lower bound in the $[2, \infty]$-regime is simply $1/\sqrt{2}$ times the upper bound. If we set $\Lambda = T^{1/p} B$, then the multi-category bound for the $[1, 2]$-regime can be compared to the one given in [15], which is larger by a factor of $O\sqrt{q}$. This improvement is however exclusively due to our trick of staying with Rademacher variables when eliminating the loss functions.

The norms in the lemma above are not very useful for multi-task learning, as the bounds show no improvement as the number of tasks increases. This is different for the trace-norm constrained class $\mathcal{W}_{\mathrm{tr}}$, for which we have the following result, which already exhibits a typical behaviour of composite classes. The proof of a more general version is given in the appendix.

Theorem 3.2

$$R_\alpha\left(\mathcal{W}_{tr}, \mathbf{x}\right) \leq B\ \theta_\alpha \left(\sqrt{\frac{2\left(\ln\left(nT\right) + 1\right)}{nT}} + \sqrt{\frac{\lambda_{\max}(\hat{C})}{n}} \right).$$

If we divide this bound by the above lower bound for regularization with the Hilbert Schmidt norm, we obtain

$$\frac{R_\alpha\left(\mathcal{W}_{\mathrm{tr}}, \mathbf{x}\right)}{R_\alpha\left(\mathcal{W}_{2,2}, \mathbf{x}\right)} \leq 2\sqrt{\frac{\ln\left(nT\right) + 1}{T}} + \sqrt{\frac{2\lambda_{\max}(\hat{C})}{\mathrm{tr}(\hat{C})}},$$

a quotient, which highlights the potential benefits of composite classes. As T increases the second term becomes dominant. The quotient $\lambda_{\max}(\hat{C})/\mathrm{tr}(\hat{C})$ can be seen as the inverse of an effective data-dimension. Indeed for whitened data $\mathrm{tr}(\hat{C}) = d\ \lambda_{\max}(\hat{C})$, if d is the number of nonzero eigenvalues of $\hat{C}$. The relative estimation benefit of the intermediate representation increases with the number T of classes or tasks and with the effective dimensionality of the data. This appears to be a rather general feature of composite vector-valued classes, also in the nonlinear case.

3.3 Composite Classes and Representation Learning

We now consider function classes $\mathcal{V} \circ \phi \circ \mathcal{W}$ of the form

$$\begin{array}{ccccccc} & & \mathcal{W} & & \phi & & \mathcal{V} & \\ x \in H & \longrightarrow & \mathbb{R}^K & \longrightarrow & H' & \rightarrow & \mathbb{R}^T . \end{array}$$

Here inputs $x \in H$ are first mapped to $\mathbb{R}^K$ by a linear function W from a class $\mathcal{W}$. The vector Wx is then mapped to another Hilbert-space H' by a fixed Lipschitz feature map $\phi : \mathbb{R}^K \to H'$. Finally $\phi(Wx)$ is mapped to the T-dimensional vector $V\phi(Wx)$ by the linear map V chosen from $\mathcal{V}$.

For $W \in \mathcal{W}$ we consider the constraints $\|W\|_{2,\infty} \leq b_\infty$, $\|W\|_{2,2} \leq b_2$ and $\|W\|_{2,1} \leq b_1$, denoting the respective classes by $\mathcal{W}_{2,\infty}$, $\mathcal{W}_{2,2}$, and $\mathcal{W}_{2,1}$. For $\mathcal{V}$ we take the constraint $\|V\|_{2,\infty} \leq a$. This choice allows us to vary T and keep a fixed at the same time. For the "activation function" ϕ we assume a Lipschitz constant L_ϕ. We make the simplifying assumption that $\phi(0) = 0$.

The function ϕ makes the model quite general. Suppose first that $H' = \mathbb{R}^K$. If ϕ is the identity function we obtain a linear class, defined through its factorization, much like the case of trace-norm regularization discussed earlier. If the components of ϕ are sigmoids or the popular rectilinear activation functions, we obtain a rather standard neural network with hidden layer, but ϕ could also include inter-unit interactions, such as poolings or lateral inhibitions (see, e.g. [10, 13]) as long as it observes the Lipschitz condition.

However, the dimension of H' need not be K and ϕ could be defined by a radial basis function network with fixed centers or it could also be the feature-map induced by some kernel on $\mathbb{R}^T$, say a Gaussian kernel of width Δ, in which case $L_\phi = 2/\Delta$. To enforce $\phi(0) = 0$ we need to translate the original feature map ψ of the Gaussian kernel as $\phi(x) = \psi(x) - \psi(0)$.

Here the underlying assumption is, that there is a common K-dimensional representation of the data in which the data has sufficient separation properties, but the separating functions may be highly nonlinear.

Theorem 3.3 *There are universal constants c_1 and c_2 such that under the above conditions*

$$\begin{aligned} R_\alpha(\mathcal{V}\phi(\mathcal{W}_{2,\infty}), \mathbf{x}) &\leq L_\phi a b_\infty \theta_\alpha \left(c_1 K \sqrt{\frac{\operatorname{tr}(\hat{C})}{nT}} + c_2 \sqrt{\frac{K \lambda_{\max}(\hat{C})}{n}} \right) \\ R_\alpha(\mathcal{V}\phi(\mathcal{W}_{2,2}), \mathbf{x}) &\leq L_\phi a b_2 \theta_\alpha \left(c_1 \sqrt{\frac{K \operatorname{tr}(\hat{C})}{nT}} + c_2 \sqrt{\frac{\lambda_{\max}\left(\hat{C}\right)}{n}} \right) \\ R_\alpha(\mathcal{V}\phi(\mathcal{W}_{2,1}), \mathbf{x}) &\leq L_\phi a b_1 \theta_\alpha \left(c_1 \sqrt{\frac{2\operatorname{tr}(\hat{C}) + 8\lambda_{\max}(\hat{C}) \ln K}{nT}} + c_2 \sqrt{\frac{\lambda_{\max}(\hat{C})}{n}} \right) . \end{aligned}$$

We highlight some implications of the above theorem.

1. The bounds differ in their dependence on the dimension K of the hidden layer which is linear, radical and logarithmic respectively. For $\mathcal{W}_{2,1}$ the dependence on K is logarithmic and scales only with $\lambda_{\max}(\hat{C})$.
2. In the case of multi-task learning with $\mathcal{W}_{2,2}$ and $\mathcal{W}_{2,1}$ the dependence on K vanishes in the limit $T \to \infty$. In this limit the first term in parenthesis vanishes in all three cases, leaving only the second term.
3. Multi-category learning requires more data with $\theta = \sqrt{T}$, but if we take a simultaneous limit in T and n such that T/n remains bounded, then the behaviour is the same as for multi-task learning with $T \to \infty$.
4. In both cases the second term becomes dominant for large T. For the first bound crudely setting $\lambda_{\max}(\hat{C}) = 1/d$ this term scales with $\sqrt{K/d}$ and exhibits the benefit of the shared representation as that of dimensional reduction. A similar interpretation holds for the other bounds with some implicit dependence of b_2 and b_1 on the dimension of the representation.

The proof uses the following recent result on the expected suprema of Gaussian processes [18]. For a set $Y \subseteq \mathbb{R}^m$ the Gaussian width $G(Y)$ is defined as

$$G(Y) = \mathbb{E} \sup_{y \in Y} \langle \gamma, y \rangle = \mathbb{E} \sup_{y \in Y} \sum_{i=1}^{m} \gamma_i y_i,$$

where $\gamma = (\gamma_1, \ldots, \gamma_m)$ is a vector of independent standard normal variables.

Theorem 3.4 *Let $Y \subseteq \mathbb{R}^n$ have (Euclidean) diameter $D(Y)$ and let $\mathcal{F}$ be a class of functions $f : Y \to \mathbb{R}^m$, all of which have Lipschitz constant at most $L(\mathcal{F})$. Let $\mathcal{F}(Y) = \{f(y) : f \in \mathcal{F}, y \in Y\}$. Then for any $y_0 \in Y$*

$$G(\mathcal{F}(Y)) \leq c_1 L(\mathcal{F}) G(Y) + c_2 D(Y) Q(\mathcal{F}) + G(\mathcal{F}(y_0)), \tag{5}$$

where c_1 and c_2 are universal constants and

$$Q(\mathcal{F}) = \sup_{\mathbf{y}, \mathbf{y}' \in Y,\ \mathbf{y} \neq \mathbf{y}'} \mathbb{E} \sup_{f \in \mathcal{F}} \frac{\langle \gamma, f(\mathbf{y}) - f(\mathbf{y}') \rangle}{\|\mathbf{y} - \mathbf{y}'\|}.$$

We refer to the appendix for statement and proof of a more general version going beyond 1-vs-all multi-category and multi-task learning.

Idea of proof for Theorem 3.3. We use Theorem 3.4 by setting

$$Y = \left\{ W\mathbf{x} = (\langle w_k, x_i \rangle)_{k \leq K,\ i \leq N} : W \in \mathcal{W} \right\} \subseteq \mathbb{R}^{KN}$$

where $\mathcal{W}$ will be either $\mathcal{W}_{2,\infty}$, $\mathcal{W}_{2,2}$ or $\mathcal{W}_{2,1}$. Note that the cardinality $|I_t|$ is either N or n in the cases considered here. For $\mathcal{F}$ we take the set of functions

$$\left\{(y_{ki}) \in \mathbb{R}^{KN} \mapsto (\langle v_t, \phi(y_i)\rangle)_{t\leq T, i\in I_t} \in \mathbb{R}^{T|I_t|} : v \in \mathcal{V}\right\}$$

restricted to Y, so $\mathcal{F}(Y)$ is a subset of $\mathbb{R}^{T^2 n}$ for multi-category and $\mathbb{R}^{Tn}$ for multi-task learning. This again accounts for the additional factor of $\sqrt{T}$ for the complexity of multi-category learning. By a well known bound on Rademacher averages in terms of Gaussian averages [14]

$$\begin{aligned}\mathbb{E}\sup_{W\in\mathcal{V}, W\in\mathcal{W}} \sum_t \sum_{i\in I_t} \epsilon_{ti} V\phi(Wx_i) &\leq \sqrt{\frac{\pi}{2}}\mathbb{E}\sup_{W\in\mathcal{V}, W\in\mathcal{W}} \sum_t \sum_{i\in I_t} \gamma_{ti} V\phi(Wx_i) \\ &= \sqrt{\frac{\pi}{2}} G(\mathcal{F}(Y)). \end{aligned} \tag{6}$$

To bound $G(\mathcal{F}(Y))$ we then just need to bound the individual components of the right hand side of equation (7), namely the largest Lipschitz constant $L(\mathcal{F})$, the differential Gaussian width $Q(\mathcal{F})$, the diameter $D(Y)$ and the Gaussian width $G(Y)$. We needn't worry about $G(\mathcal{F}(y_0))$, because we are free to choose y_0, so we can set it to 0. Then $f(0) = 0$ for all $f \in \mathcal{F}$, whence $G(\mathcal{F}(y_0)) = 0$. For the bounds on $L(\mathcal{F})$, $Q(\mathcal{F})$, $D(Y)$ and $G(Y)$ we refer to the appendix. □

4 Conclusion

We presented a framework to derive Rademacher bounds for a wide class of vector-valued functions combined with Lipschitz losses. We studied in parallel the case of multi-task and multi-category learning. To our knowledge our framework allows to derive bounds for more general classes of vector-valued function and loss functions than currently possible, while still improving over existing bounds [15, 17] in special cases. In particular, we illustrate how bounds can be derived for neural networks with one hidden layer and rather general nonlinear activation functions.

In the future, it would be valuable to study more examples of the loss functions included in the setting. In addition to one-vs-one classification, which we briefly mentioned in the paper, these could include multi-label classification or hybrid multi-task learning, in which each task is itself a multi-category or multi-label problem. Another interesting direction of research is to extend our analysis to neural networks with more than one hidden layer. Although the proof technique presented in Section 3.3 could naturally be extended to derive such bounds, it seems important to study improvement in the large constants appearing in Theorem 3.4 (see [18]) in order to avoid explosion of the constants in bounds for deep networks.

A Appendix

For the convenience of the reader we restate in greater generality the results contained in the main body of the paper. The ϵ_i or ϵ_{ti} are throughout independent Rademacher variables.

A.1 Mixed Norms

In this section we prove a more general result implying Theorem 3.1.

Theorem A.1 *We have that:*

(i) For $p \in [2, \infty]$

$$\frac{B}{\sqrt{2}N} \sum_{t=1}^{T} \sqrt{|I_t| \operatorname{tr}(\hat{C}_t)} \leq R_I\left(\mathcal{W}_{2,\infty}, \mathbf{x}\right) \leq R_I\left(\mathcal{W}_{2,p}, \mathbf{x}\right) \leq \frac{B\sqrt{T}}{N} \sqrt{\sum_{t=1}^{T} |I_t| \operatorname{tr}(\hat{C}_t)}.$$

(ii) For $p \in [1, 2]$ *and* $1/p + 1/q = 1$ *if* $\sum_{i \in I_t} \|x_i\|^2 \geq q^{-1}$ *then*

$$R_I\left(\mathcal{W}_{2,2}, \mathbf{x}\right) \leq R_I(\mathcal{W}_{2,p}, \mathbf{x}) \leq \frac{T^{1/p} B\sqrt{q}}{N} \left(2 \sum_t \sqrt{|I_t| \operatorname{tr}(\hat{C}_t)}^q\right)^{1/q},$$

where $1/p + 1/q = 1$.

(iii) For 1-vs-all multi-category learning the condition $\sum_{i \in I_t} \|x_i\|^2 \geq q^{-1}$ *can be omitted and the bound in (ii) can be simplified to*

$$R_{I^{\mathrm{mc}}}\left(\mathcal{W}_{2,p}, \mathbf{x}\right) \leq B \sqrt{\frac{qT \operatorname{tr}(\hat{C})}{n}}.$$

Proof. (i) We have

$$\begin{aligned}
\frac{B}{\sqrt{2}} \sum_t \sqrt{|I_t| \operatorname{tr}(\hat{C}_t)} &= \frac{B}{\sqrt{2}} \sum_t \sqrt{\sum_{i \in I_t} \|x_i\|^2} = \frac{B}{\sqrt{2}} \sum_t \sqrt{\mathbb{E} \left\| \sum_{i \in I_t} \epsilon_i x_i \right\|^2} \\
&\leq B \sum_t \mathbb{E} \left\| \sum_{i \in I_t} \epsilon_i x_i \right\| = \sum_t \mathbb{E} \sup_{w, \|w\| \leq B} \left\langle w, \sum_{i \in I_t} \epsilon_i x_i \right\rangle \\
&= N\, R_I\left(\mathcal{W}_{2,\infty}\right) \leq N\, R_I\left(\mathcal{W}_{2,p}\right) \leq N\, R_I\left(\mathcal{W}_{2,2}\right) \\
&= \mathbb{E} \sup_{W \in \mathcal{W}_{2,2}} \sum_t \left\langle w_t, \sum_{i \in I_t} \epsilon_i x_i \right\rangle = B\sqrt{T} \mathbb{E} \sqrt{\sum_t \left\| \sum_{i \in I_t} \epsilon_i x_i \right\|^2} \\
&\leq B\sqrt{T} \sqrt{\sum_t \sum_{i \in I_t} \|x_i\|^2} = B\sqrt{T} \sqrt{\sum_t |I_t| \operatorname{tr}(\hat{C}_t)}
\end{aligned}$$

where we used Szarek's inequality (Theorem 5.20 [6]) in the first inequality. The next inequalities follow from $\mathcal{W}_{2,\infty} \subseteq \mathcal{W}_{2,p} \subseteq \mathcal{W}_{2,2}$. For the last inequality we use Jensen's.

(ii) The first inequality is $\mathcal{W}_{2,2} \subseteq \mathcal{W}_{2,p}$. Then let $X_t = \left\|\sum_{i\in I_t} \epsilon_i x_i\right\|$, so that $\mathbb{E}X_t \leq \sqrt{\sum_{i\in I_t} \|x_i\|^2}$. By the bounded difference inequality (see [6]) for $s \geq 0$

$$\Pr\left\{X_t > \mathbb{E}X_t + s\right\} \leq \exp\left(\frac{-s^2}{2\sum_{i\in I_t}\|x_i\|^2}\right),$$

so with integration by parts

$$\begin{aligned}
\mathbb{E}\left[X_t^q\right] &\leq \mathbb{E}X_t + q\int_0^\infty s^{q-1}\Pr\left\{X > \mathbb{E}X + s\right\} ds^q \\
&\leq \mathbb{E}X_t + q\int_0^\infty s^{q-1}\exp\left(\frac{-s^2}{2\sum_{i\in I_t}\|x_i\|^2}\right) ds \\
&= \mathbb{E}X_t + \left(\sum_{i\in I_t}\|x_i\|^2\right)^{q/2}\left(q\int_0^\infty s^{q-1}\exp\left(\frac{-s^2}{2}\right) ds\right) \\
&\leq \left(\sum_{i\in I_t}\|x_i\|^2\right)^{1/2} + \left(q\sum_{i\in I_t}\|x_i\|^2\right)^{q/2} \leq 2\left(q\sum_{i\in I_t}\|x_i\|^2\right)^{q/2},
\end{aligned}$$

where the third inequality follows from a comparison of the integral with the moments of the standard normal distribution, and the last follows from $\sum_{i\in I_t}\|x_i\|^2 \geq q^{-1}$. Thus

$$\begin{aligned}
R_I\left(\mathcal{W}_{2,p}\right) &= \frac{1}{N}\mathbb{E}\sup_{\|W\|_{2,p}\leq T^{1/p}B}\sum_t\sum_{i\in I_t}\langle w_t, x_i\rangle = \frac{T^{1/p}B}{N}\mathbb{E}\left(\sum_t X_t^q\right)^{1/q} \\
&\leq \frac{T^{1/p}B}{N}\left(\sum_t \mathbb{E}X_t^q\right)^{1/q} \leq \frac{T^{1/p}B\sqrt{q}}{N}\left(2\sum_t\left(\sum_{i\in I_t}\|x_i\|^2\right)^{q/2}\right)^{1/q} \\
&= \frac{2^{1/q}T^{1/p}B\sqrt{q}}{N}\left(\sum_t\left(|I_t|\operatorname{tr}(\hat{C}_t)\right)^{q/2}\right)^{1/q}.
\end{aligned}$$

(iii) The case of 1-vs-all multi-category learning is simpler because $I_t = \{1,\ldots,N\}$ and we can interchange summation over t and i. Then we can essentially proceed as in[15] and use the $1/q$-strong convexity of $\frac{1}{2}\|W\|_{2,p}^2$ w.r.t. $\|W\|_{2,p}$. In Corollary 4 of [11] let $\lambda > 0$ and $u = W$ and $v_i = \lambda\left(\epsilon_{1i}x_i,\ldots,\epsilon_{Ti}x_i\right)$ and use $\frac{1}{2}\|W\|_{2,p}^2 \leq \frac{1}{2}\left(T^{1/p}B\right)^2 = f_{\max}(u)$ to obtain

$$\sum_{i=1}^N \langle W, \lambda\left(\epsilon_{1i}x_i,...,\epsilon_{Ti}x_i\right)\rangle_2 \leq \sum_{i=1}^N \langle \nabla f\left(v_{1:i-1}\right), v_i\rangle + \frac{1}{2}\left(T^{1/p}B\right)^2 + \frac{q\lambda^2}{2}\sum_{i=1}^N \|(\epsilon_{1i}x_i,...,\epsilon_{Ti}x_i)\|_{2q}^2,$$

where $\langle \cdot, \cdot \rangle_2$ is the Hilbert-Schmidt inner product. Take the supremum in W and then the expectation. The first term on the r.h.s. above vanishes. Dividing by λ and optimizing in λ gives

$$\mathbb{E} \sup_W \sum_{i=1}^n \langle W, \epsilon_{1i} x_i, \ldots, \epsilon_{Ti} x_i \rangle \leq \left(T^{1/p} B \right) \sqrt{q \sum_{i=1}^n \mathbb{E} \left\| (\epsilon_{1i} x_i, \ldots, \epsilon_{Ti} x_i) \right\|_{2q}^2}.$$

Now

$$\mathbb{E} \left\| (\epsilon_{1i} x_i, \ldots, \epsilon_{Ti} x_i) \right\|_{2q}^2 = \mathbb{E} \left(\sum_t \left\| \epsilon_{ti} x_i \right\|^q \right)^{2/q} \leq T^{2/q} \left\| x_i \right\|^2$$

so

$$R_{I^{\text{mc}}} (\mathcal{W}_{2,p}) = \frac{1}{N} \mathbb{E} \sum_{i=1}^N \langle W, \epsilon_{1i} x_i, \ldots, \epsilon_{Ti} x_i \rangle \leq \frac{TB}{N} \sqrt{q \sum_{i=1}^N \left\| x_i \right\|^2} = B \sqrt{\frac{qT \operatorname{tr}(\hat{C})}{n}}.$$

□

Note that the (very harmless) condition $\sum_{i \in I_t} \left\| x_i \right\|^2 \geq q^{-1}$ in part (iii) is automatically satisfied if $\left\| x_i \right\| = 1$.

A.2 Trace Norm Constraints

In this section we prove the following result, which contains Theorem 3.2 as as special case and improves over [17] which only applies to the multi-task learning setting.

Theorem A.2

$$R_I (\mathcal{W}_{tr}, \mathbf{x}) \leq \frac{B}{N} \sqrt{2T \max_t |I_t| \operatorname{tr}(\hat{C}_t) (\ln N + 1)} + \frac{B}{N} \sqrt{T \lambda_{\max} \left(\sum_t |I_t| \hat{C}_t \right)}.$$

For the proof we use $\left\| . \right\|_\infty$ to denote the operator norm on H and $\succeq$ and $\preceq$ to refer to the ordering induced by the cone of positive operators. For $x \in H$ we define the rank-1 operator Q_x on H by $Q_x v = \langle v, x \rangle x$. We use the following result, the proof of which can be found in [17].

Theorem A.3 *Let $M \subseteq H$ be a subspace of dimension d and suppose that $A_1, \ldots, A_N$ are independent random operators satisfying $A_k \succeq 0$, $Ran(A_k) \subseteq M$ a.s. and*

$$\mathbb{E} A_k^m \preceq m! R^{m-1} \mathbb{E} A_k$$

for some $R \geq 0$, all $m \in \mathbb{N}$ and all $k \in \{1, \ldots, N\}$. Then

$$\sqrt{\mathbb{E} \left\| \sum_k A_k \right\|_\infty} \leq \sqrt{\left\| \mathbb{E} \sum_k A_k \right\|_\infty} + \sqrt{R (\ln \dim (M) + 1)}.$$

Lemma A.4 *Let $x_1, \ldots, x_n$ be in $\mathbb{R}^d$ and denote*

$$\alpha = \sum_{i=1}^{n} \|x_i\|^2 .$$

Define a random vector by $V = \sum_i \epsilon_i x_i$. Then for $p \geq 1$

$$\mathbb{E}\left[Q_V^p\right] \preceq (2p-1)!!\alpha^{p-1}\mathbb{E}\left[Q_V\right],$$

where $(2p-1)!! = \prod_{i=1}^{p}(2i-1) = (2p-1)\left(2\left(p-1\right)-1\right) \times \cdots \times 5 \times 3 \times 1$.

Proof. Let $v \in \mathbb{R}^d$ be arbitrary. By the definition of V and Q_V we have for any $v \in \mathbb{R}^d$ that

$$\left\langle \mathbb{E}\left[Q_V^p\right] v, v \right\rangle = \sum_{j_1,\ldots,j_{2p}=1}^{n} \mathbb{E}\left[\epsilon_{j_1}\epsilon_{j_2}\cdots\epsilon_{j_{2p}}\right] \left\langle v, x_{j_1}\right\rangle \left\langle x_{j_1}, x_{j_2}\right\rangle \cdots \left\langle x_{j_{2p}}, v\right\rangle .$$

The properties of independent Rademacher variables imply that $\mathbb{E}\left[\epsilon_{i_1}\epsilon_{i_2}\cdots\epsilon_{i_{2p}}\right] = 0$ unless the sequence $\mathbf{i} = (i_1, \ldots, i_{2p})$ has the property that each index i_k occurs in it an even number of times, in which case $\mathbb{E}\left[\epsilon_{i_1}\epsilon_{i_2}\cdots\epsilon_{i_{2p}}\right] = 1$. Let us call sequences with this property admissible. Thus

$$\begin{aligned}
\left\langle E\left[Q_w^p\right] v, v\right\rangle &= \sum_{\mathbf{i}\text{ admissible}} \left\langle v, x_{i_1}\right\rangle \left\langle x_{i_2}, x_{i_3}\right\rangle \cdots \left\langle x_{i_{2p}}, v\right\rangle \\
&\leq \sum_{\mathbf{i}\text{ admissible}} \left|\left\langle v, x_{i_1}\right\rangle\right| \prod_{k=2}^{2p-1} \|x_{i_k}\| \left|\left\langle x_{i_{2p}}, v\right\rangle\right|,
\end{aligned}$$

using Cauchy-Schwarz. For every admissible sequence $\mathbf{i}$ there exists at least one partition π of $\{1, .., 2p\}$ into p pairs (l, r) with $l < r$, such that the indices i_{k_1} and i_{k_2} are equal, whenever k_1 and k_2 belong to the same pair. Let us denote the latter condition by $\mathbf{i} \sim \pi$. It is easy to show by induction that there are $(2p-1)!!$ such partitions into pairs. Given π we can write $\{1, \ldots, 2p\} = L_\pi \cup R_\pi$, where $L_\pi = \{l : \exists\,(l, r) \in \pi\}$ and $R_\pi = \{r : \exists\,(l, r) \in \pi\}$. We always have $1 \in L_\pi$ and $2p \in R_\pi$ and $|L_\pi| = |R_\pi| = p$. Thus

$$\begin{aligned}
\left\langle E\left[Q_w^p\right] v, v\right\rangle &\leq \sum_{\pi}\sum_{\mathbf{i}\sim\pi} \left|\left\langle v, x_{i_1}\right\rangle\right| \prod_{k=2}^{2p-1} \|x_{i_k}\| \left|\left\langle x_{i_{2p}}, v\right\rangle\right| \\
&= \sum_{\pi}\sum_{\mathbf{i}\sim\pi} \left(\left|\left\langle v, x_{i_1}\right\rangle\right| \prod_{k=2, i_k\in L_\pi}^{2p-1} \|x_{i_k}\|\right) \left(\left|\left\langle x_{i_{2p}}, v\right\rangle\right| \prod_{k=2, i_k\in R_\pi}^{2p-1} \|x_{i_k}\|\right) \\
&\leq \sum_{\pi}\sum_{\mathbf{i}\sim\pi} \left\langle v, x_{i_1}\right\rangle^2 \prod_{k=2, i_k\in L_\pi}^{2p-1} \|x_{i_k}\|^2 .
\end{aligned}$$

The last step follows from the Cauchy-Schwarz inequality and realizing that the two resulting factors are equal by symmetry. But for $\mathbf{i} \sim \pi$ we just need to sum over the indices in L_π, the others

being constrained to be equal. Thus, writing $L_\pi = \{l_1, \ldots, l_p\}$ such that $l_1 = 1$ the last expression above is just

$$\begin{aligned}
&\sum_\pi \sum_{i_1,\ldots,i_p} \langle v, x_{i_1}\rangle^2 \prod_{k=2}^{p} \|x_{i_k}\|^2 \\
= \;& (2p-1)!! \left(\sum_{i=1}^n \|x_i\|^2\right)^{p-1} \left\langle \sum_{i=1}^n Q_{x_i} v, v \right\rangle \\
= \;& (2p-1)!! \left(\sum_{i=1}^n \|x_i\|^2\right)^{p-1} \langle \mathbb{E}\left[Q_V\right] v, v\rangle \; .
\end{aligned}$$

The conclusion follows since for symmetric matrices $(\forall v, \langle Av, v\rangle \leq \langle B, v, v\rangle) \implies A \preceq B$. □

Proof of Theorem A.2. We have

$$R_I\left(\mathcal{W}_{tr}, \mathbf{x}\right) = \frac{1}{N}\mathbb{E} \sup_{W\in\mathcal{W}_{\mathrm{tr}}} \sum_t \sum_{i\in I_t} \epsilon_{ti} \langle w_t, x_i\rangle = \frac{1}{N}\mathbb{E} \sup_{W\in\mathcal{W}_{\mathrm{tr}}} \mathrm{tr}(W^* D),$$

where the random operator $D : H \to \mathbb{R}^T$ is defined for $v \in H$ by $(Dv)_t = \left\langle v, \sum_{i\in I_t} \epsilon_{ti} x_i\right\rangle$. Hölder's inequality gives

$$R_I\left(\mathcal{W}_{\mathrm{tr}}, \mathbf{x}\right) \leq \frac{B\sqrt{T}}{N} \mathbb{E}\left\|D\right\|_\infty .$$

We proceed to bound $\mathbb{E}\left\|D\right\|_\infty$. Let V_t be the random vector $V_t = \sum_{i\in I_t}^{n_t} \epsilon_{ti} x_i$ and recall that the corresponding rank-one operator Q_{V_t} is defined by $Q_{V_t} v = \langle v, V_t\rangle V_t = \left\langle v, \sum_{i\in I_t}^{n_t} \epsilon_{ti} x_i\right\rangle \sum_{i\in I_t}^{n_t} \epsilon_{ti} x_i$. Then $D^* D = \sum_{t=1}^T Q_{V_t}$, so by Jensen's inequality

$$\mathbb{E}\left\|D\right\|_\infty \leq \sqrt{\mathbb{E}\left\|\sum_t Q_{V_t}\right\|_\infty} \; .$$

The range of any of the realizations of Q_{V_t} lies in the span of the x_i which has less than N. By Lemma A.4 we have with $\alpha_t = \sum_{i\in I_t} \|x_i\|^2$

$$\mathbb{E}\left[(Q_{Vt})^m\right] \preceq (2p-1)!! \alpha_t^{m-1} \mathbb{E}\left[Q_{V_t}\right] \preceq m! \left(2\max_t \alpha_t\right)^{m-1} \mathbb{E}\left[Q_{V_t}\right],$$

so Theorem A.3 with $R = 2\max_t \alpha_t$ and $d = N$ now gives

$$\sqrt{\mathbb{E}\left\|\sum_t Q_{V_t}\right\|_\infty} \leq \sqrt{2\max_t \alpha_t \left(\ln N + 1\right)} + \sqrt{\left\|\mathbb{E}\sum_t Q_{V_t}\right\|_\infty} \; .$$

But $\mathbb{E}[Q_{V_t}] = \sum_{i \in I_t} Q_{x_i} = |I_t| \hat{C}_t$, so

$$\begin{aligned} R_I(\mathcal{W}_{\mathrm{tr}}, \mathbf{x}) &\leq \frac{B\sqrt{T}}{N} \mathbb{E} \|D\|_\infty \leq \frac{B\sqrt{T}}{N} \sqrt{\mathbb{E} \left\| \sum_t Q_{V_t} \right\|_\infty} \\ &\leq \frac{B}{N} \sqrt{2T \max_t |I_t| \operatorname{tr}(\hat{C}_t) (\ln N + 1)} + \sqrt{T \left\| \sum_t |I_t| \hat{C}_t \right\|_\infty}. \end{aligned}$$

□

A.3 Nonlinear Compositions

For the statement of a general version of Theorem 3.3 we extend the definition of θ_{mc} and θ_{mt} by setting for any map $I : \{1, \ldots, T\} \to 2^{\{1,\ldots,N\}}$

$$\theta_I = \inf \left\{ \theta : \forall (a_1, \ldots, a_N), a_i \geq 0, \ \sum_{t=1}^{T} \sum_{i \in I_t} a_i \leq \theta^2 \sum_{i=1}^{N} a_i \right\}.$$

This definition coincides with the previous one in the case of multi-task and 1-vs-all multi-category learning.

Theorem A.5 *There are universal constants c_1 and c_2 such that under the above conditions*

$$\begin{aligned} R_I(\mathcal{V}\phi(\mathcal{W}_{2,\infty}), \mathbf{x}) &\leq L_\phi a b_\infty \theta_I \left(c_1 K \sqrt{\frac{\operatorname{tr}(\hat{C})}{nT}} + c_2 \sqrt{\frac{K \lambda_{\max}(\hat{C})}{n}} \right) \\ R_I(\mathcal{V}\phi(\mathcal{W}_{2,2}), \mathbf{x}) &\leq L_\phi a b_2 \theta_I \left(c_1 \sqrt{\frac{K \operatorname{tr}(\hat{C})}{nT}} + c_2 \sqrt{\frac{\lambda_{\max}\left(\hat{C}\right)}{n}} \right) \\ R_I(\mathcal{V}\phi(\mathcal{W}_{2,1}), \mathbf{x}) &\leq L_\phi a b_1 \theta_I \left(c_1 \sqrt{\frac{2\operatorname{tr}(\hat{C}) + 8\lambda_{\max}(\hat{C}) \ln K}{nT}} + c_2 \sqrt{\frac{\lambda_{\max}(\hat{C})}{n}} \right). \end{aligned}$$

The proof uses the following recent result on the expected suprema of Gaussian processes [18]. For a set $Y \subseteq \mathbb{R}^m$ the Gaussian width $G(Y)$ is defined as

$$G(Y) = \mathbb{E} \sup_{y \in Y} \langle \gamma, y \rangle = \mathbb{E} \sup_{y \in Y} \sum_{i=1}^{m} \gamma_i y_i,$$

where $\gamma = (\gamma_1, \ldots, \gamma_m)$ is a vector of independent standard normal variables.

Theorem A.6 *Let $Y \subseteq \mathbb{R}^n$ have (Euclidean) diameter $D(Y)$ and let $\mathcal{F}$ be a class of functions $f : Y \to \mathbb{R}^m$, all of which have Lipschitz constant at most $L(\mathcal{F})$. Let $\mathcal{F}(Y) = \{f(y) : f \in \mathcal{F}, y \in Y\}$. Then for any $y_0 \in Y$*

$$G(\mathcal{F}(Y)) \leq c_1 L(\mathcal{F}) G(Y) + c_2 W(Y) Q(\mathcal{F}) + G(\mathcal{F}(y_0)), \tag{7}$$

where c_1 and c_2 are universal constants and

$$Q(\mathcal{F}) = \sup_{\mathbf{y},\mathbf{y}' \in Y,\, \mathbf{y} \neq \mathbf{y}'} \mathbb{E} \sup_{f \in \mathcal{F}} \frac{\langle \gamma, f(\mathbf{y}) - f(\mathbf{y}') \rangle}{\|\mathbf{y} - \mathbf{y}'\|}.$$

Proof of Theorem 3.3. We will use Theorem A.6 by setting

$$Y = \left\{ W\mathbf{x} = (\langle w_k, x_i \rangle)_{k \leq K,\, i \leq N} : W \in \mathcal{W} \right\} \subseteq \mathbb{R}^{KN}$$

where $\mathcal{W}$ will be either $\mathcal{W}_{2,\infty}$, $\mathcal{W}_{2,2}$ or $\mathcal{W}_{2,1}$. For $\mathcal{F}$ we take the set of functions

$$\left\{ (y_{ki}) \in \mathbb{R}^{KN} \mapsto (\langle v_t, \phi(y_i) \rangle)_{t \leq T, i \in I_t} \in \prod_{t=1}^{T} \mathbb{R}^{|I_t|} : v \in \mathcal{V} \right\}$$

restricted to Y. By a well known bound on Rademacher averages in terms of Gaussian averages [14]

$$\begin{aligned} \mathbb{E} \sup_{W \in \mathcal{V}, W \in \mathcal{W}} \sum_t \sum_{i \in I_t} \epsilon_{ti} V \phi(W x_i) &\leq \sqrt{\frac{\pi}{2}} \mathbb{E} \sup_{W \in \mathcal{V}, W \in \mathcal{W}} \sum_t \sum_{i \in I_t} \gamma_{ti} V \phi(W x_i) \\ &= \sqrt{\frac{\pi}{2}} G(\mathcal{F}(Y)). \end{aligned} \tag{8}$$

To bound $G(\mathcal{F}(Y))$ we then just need to bound the terms in the right hand side of equation (7)

Since $\phi(0) = 0$, we can at once set $G(\mathcal{F}(y_0)) = 0$, by setting $0 = y_0$, so $f(0) = 0$ for all $f \in \mathcal{F}$.

Bounding the Lipschitz constant. For any $v \in \mathcal{V}$ and $y, y' \in Y \subseteq \mathbb{R}^{KN}$,

$$\begin{aligned} \sum_{t, i \in I_t} \left(\langle v_t, \phi(y_i) \rangle - \langle v_t, \phi(y_i') \rangle \right)^2 &\leq \sum_t \|v_t\|^2 \sum_{i \in I_t} \left\| \phi(y_i) - \phi(y_i') \right\|^2 \\ &\leq a^2 L_\phi^2 \sum_t \sum_{i \in I_t} \left\| y_i - y_i' \right\|^2 \leq a^2 L_\phi^2 \theta_I^2 \left\| y - y' \right\|^2, \end{aligned}$$

so $L(\mathcal{F}) \leq a L_\phi \theta_I$.

Bounding $Q(\mathcal{F})$. Again with $y, y' \in Y$

$$\begin{aligned}
&\mathbb{E}\sup_{f\in\mathcal{F}}\left\langle\gamma, f(\mathbf{y}) - f(\mathbf{y}')\right\rangle \\
&= \mathbb{E}\sup_{v\in\mathcal{V}}\sum_{ti}\gamma_{ti}\left(\langle v_t, \phi(y_i)\rangle - \langle v_t, \phi(y_i')\rangle\right) = \mathbb{E}\sup_{v\in\mathcal{V}}\sum_t\left\langle v_t, \sum_{i\in I_t}\gamma_{ti}\left(\phi(y_i) - \phi(y_i')\right)\right\rangle \\
&\le a\mathbb{E}\sum_t\left\|\sum_{i\in I_t}\gamma_{ti}\left(\phi(y_i) - \phi(y_i')\right)\right\| \le \sqrt{T}a\left(\sum_t E\left\|\sum_{i\in I_t}\gamma_{ti}\left(\phi(y_i) - \phi(y_i')\right)\right\|^2\right)^{1/2} \\
&\le aL_\phi\sqrt{T}\left(\sum_t\sum_{i\in I_t}\|y_i - y_i'\|^2\right)^{1/2} \le aL_\phi\theta_I\sqrt{T}\|y - y'\|,
\end{aligned}$$

so $Q(\mathcal{F}) \le aL_\phi\theta_I\sqrt{T}$.

Bounding the diameters. We have

$$\begin{aligned}
D(\mathcal{W}\mathbf{x}) &\le 2\sqrt{\sup_W\sum_{ki}\langle w_k, x_i\rangle^2} = \sqrt{\sup_W\sum_k\|w_k\|^2\sum_i\left\langle\frac{w_k}{\|w_k\|}, x_i\right\rangle^2} \\
&\le \sqrt{\sup_W\sum_k\|w_k\|^2 N\lambda_{\max}(\hat{C})} = \|W\|_{2,2}\sqrt{N\lambda_{\max}(\hat{C})}.
\end{aligned}$$

From $\|W\|_{2,2} \le \|W\|_{2,1}$ and $\|W\|_{2,2} \le \sqrt{K}\|W\|_{2,\infty}$ we obtain

$$D(\mathcal{W}_{2,\infty}) \le b_\infty\sqrt{KN\lambda_{\max}(\hat{C})}, \text{ and both } D(\mathcal{W}_{2,2}), D(\mathcal{W}_{2,1}) \le b_2\sqrt{N\lambda_{\max}(\hat{C})}.$$

Bounding the Gaussian width.

$$G(\mathcal{W}_{2,\infty}\mathbf{x}) = \mathbb{E}\sup_{W\in\mathcal{W}_\infty}\sum_k\left\langle w_k, \sum_{i\le N}\gamma_{ki}x_i\right\rangle = b_\infty\sum_k\mathbb{E}\left\|\sum_{i\le N}\gamma_{ki}x_i\right\| \le b_\infty K\sqrt{N\mathrm{tr}(\hat{C})}.$$

similarly

$$G(\mathcal{W}_{2,2}\mathbf{x}) = \mathbb{E}\sup_{W\in\mathcal{W}_2}\sum_k\left\langle w_k, \sum_{i\le N}\gamma_{ki}x_i\right\rangle = b_2\sqrt{\sum_k\mathbb{E}\left\|\sum_{i\le N}\gamma_{ki}x_i\right\|^2} \le b_\infty\sqrt{KN\ \mathrm{tr}(\hat{C})}.$$

The Gaussian width of $\mathcal{W}_1\mathbf{x}$ is a little more complicated. Let $\mathcal{W}_1^{(k)}$ be the class of linear transformations $\mathcal{W}_1^{(k)} = \{x \mapsto (0, \ldots, \langle w, x\rangle, \ldots, 0) : \|w\| \le b_1\}$, where only the k-th coordinate is different

from zero. Then $\mathcal{W}_1\mathbf{x}$ is the convex hull of $\mathcal{W}_1^{(1)}\mathbf{x} \cup \cdots \cup \mathcal{W}_1^{(K)}\mathbf{x}$. It follows from Lemma 2 in [19] that

$$\begin{aligned} G\left(\mathcal{W}_{2,1}\mathbf{x}\right) &\leq \max_k G\left(\mathcal{W}_1^{(k)}\mathbf{x}\right) + 2\sqrt{\sum_{k,i} \langle w_k, x_i\rangle^2 \ln K} \\ &\leq b_1\sqrt{N \operatorname{tr}(\hat{C})} + 2\sqrt{\sum_k \|w_k\|^2 \sum_i \left\langle \frac{w_k}{\|w_k\|}, x_i \right\rangle^2 \ln K} \\ &\leq b_1\sqrt{N \operatorname{tr}(\hat{C})} + 2b_1\sqrt{N\lambda_{\max}(\hat{C}) \ln K} \\ &\leq b_1\sqrt{2N\ \left(\operatorname{tr}(\hat{C}) + 8\lambda_{\max}\left(\hat{C}\right) \ln K\right)}. \end{aligned}$$

Collecting these bounds in Theorem A.6 and using (8) gives the three inequalities of Theorem 3.3.

□

References

[1] Y. Amit, M. Fink, N. Srebro, and S. Ullman. Uncovering shared structures in multiclass classification. In *Proceedings of the 24th international conference on Machine learning*, pages 17–24, 2007.

[2] R. K. Ando and T. Zhang. A framework for learning predictive structures from multiple tasks and unlabeled data. *Journal of Machine Learning Research*, 6, 1817–1853, 2005.

[3] M. Anthony and P. L. Bartlett. *Neural network learning: Theoretical foundations*. Cambridge University Press, 1999.

[4] P. L. Bartlett and S. Mendelson. Rademacher and Gaussian Complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.

[5] J. Baxter. A model of inductive bias learning. *Journal of Artificial Intelligence Research*, 12:149–198, 2000.

[6] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities*, Oxford University Press, 2013

[7] G. Cavallanti, N. Cesa-Bianchi, and C. Gentile. Linear algorithms for online multitask classification. *Journal of Machine Learning Research*, 11:2597–2630, 2010.

[8] K. Crammer and Y. Singer. On the algorithmic implementation of multiclass kernel-based vector machines. *Journal of Machine Learning Research*, 2, 265–292, 2002

[9] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the 2014 Conference on Computer Vision and Pattern Recognition*, pages 580–587, 2014.

[10] S. Haykin. *Neural Networks: A Comprehensive Foundation*. Prentice-Hall, 1999.

[11] S. M. Kakade, S. Shalev-Shwartz, A. Tewari. Regularization techniques for learning with matrices. *Journal of Machine Learning Research* 13:1865–1890, 2012.

[12] V. Koltchinskii and D. Panchenko. Empirical margin distributions and bounding the generalization error of combined classifiers. *Annals of Statistics*, 30(1):1–50, 2002.

[13] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE,* 86(11):2278–2324, 1998

[14] M. Ledoux, M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes.* Springer, Berlin, 1991.

[15] Y. Lei, U., Dogan, A. Binder, and M. Kloft. Multi-class SVMs: From tighter data-dependent generalization bounds to novel algorithms. In *Advances in Neural Information Processing Systems*, pages 2026–2034, 2015.

[16] A. Maurer. Bounds for linear multi-task learning. *Journal of Machine Learning Research*, 7:117–139, 2006.

[17] A. Maurer, and M. Pontil. Excess risk bounds for multitask learning with trace norm regularization. In *Proceeding of the 26th Annual Conference on Learning Theory*, pages 55–76, 2013.

[18] A. Maurer. A chain rule for the expected suprema of Gaussian processes. In *Proceedings of the 25th International Conference on Algorithmic Learning Theory*, pages 245–259, 2014

[19] A. Maurer, M. Pontil, and B. Romera-Paredes. An inequality with applications to structured sparsity and multitask dictionary learning. In *Proceedings of the 27th Conference on Learning Theory*, pages 440–460, 2014.

[20] A. Maurer, M. Pontil, and B. Romera-Paredes. The benefit of multitask representation learning. arXiv preprint arXiv:1505.06279.

[21] A. Maurer. A vector-contraction inequality for Rademacher complexities. arXiv preprint arXiv:1605.00251.

[22] R. Meir and T. Zhang. Generalization error bounds for Bayesian mixture algorithms. *Journal of Machine Learning Research*, 4:839–860, 2003.

[23] Y. Mroueh, T., Poggio, R. Rosasco, and J. Slotine. Multiclass learning with simplex coding. In *Advances in Neural Information Processing Systems*, pages 2789–2797, 2012.

[24] B. Neyshabur, R. Tomioka, and N. Srebro. Norm-based capacity control in neural networks. In *Proceedings of the 28th Conference on Learning Theory*, pages 1376–1401, 2015.

[25] D. Slepian. The one-sided barrier problem for Gaussian noise. *Bell System Tech. J.*, 41:463–501, 1962.

[26] N. Srebro and A. Shraibman. Rank, trace-norm and max-norm. In *Proceedings of the 18th Annual Conference on Learning Theory*, pages 545–560, 2005.